

# Exhibit B

[Get Started for Free](#)[Contact Us](#)

## Amazon Rekognition ▾

[Overview](#)[Use cases ▾](#)[Features ▾](#)[Pricing](#)[Resources](#)[FAQs](#)[Customers](#)[Products](#) / [Machine Learning](#) / [Amazon Rekognition](#)

# Amazon Rekognition FAQs

## General

### Q: What is Amazon Rekognition?

Amazon Rekognition is a service that makes it easy to add powerful visual analysis to your applications. Rekognition Image lets you easily build powerful applications to search, verify, and organize millions of images. Rekognition Video lets you extract motion-based context from stored or live stream videos and helps you analyze them.

Rekognition Image is an image recognition service that detects objects, scenes, activities, landmarks, faces, dominant colors, and image quality. Rekognition Image also extracts text, recognizes celebrities, and identifies inappropriate content in images. It also allows you to search and compare faces.

Rekognition Video is a video recognition service that detects activities, understands the movement of people in frame, and recognizes objects, celebrities, and inappropriate content in videos stored in Amazon S3 and live video streams. Rekognition Video detects persons and tracks them through the video even when their faces are not visible, or as the whole person might go in and out of the scene. For example, this could be used in an application that sends a real-time notification when someone delivers a package to your door. Rekognition Video allows you also to index metadata like objects, activities, scene, landmarks, celebrities, and faces that make video search easy.

### **Q: What is deep learning?**

Deep learning is a sub-field of Machine Learning and a significant branch of Artificial Intelligence. It aims to infer high-level abstractions from raw data by using a deep graph with multiple processing layers composed of multiple linear and non-linear transformations. Deep learning is loosely based on models of information processing and communication in the brain. Deep learning replaces handcrafted features with ones learned from very large amounts of annotated data. Learning occurs by iteratively estimating hundreds of thousands of parameters in the deep graph with efficient algorithms.

Several deep learning architectures such as convolutional deep neural networks (CNNs), and recurrent neural networks have been applied to computer vision, speech recognition, natural language processing, and audio recognition to produce state-of-the-art results on various tasks.

Amazon Rekognition is a part of the Amazon AI family of services. Amazon AI services use deep learning to understand images, turn text into lifelike speech, and build intuitive conversational text and speech interfaces.

### **Q: Do I need any deep learning expertise to use Amazon Rekognition?**

No. With Amazon Rekognition, you don't have to build, maintain or upgrade deep learning pipelines.

To achieve accurate results on complex computer vision tasks such as object and scene detection, face analysis, and face recognition, deep learning systems need to be tuned properly and trained with massive amounts of labeled ground truth data. Sourcing, cleaning, and labeling data

accurately is a time-consuming and expensive task. Moreover, training a deep neural network is computationally expensive and often requires custom hardware built using Graphics Processing Units (GPU).

Amazon Rekognition is fully managed and comes pre-trained for image and video recognition tasks, so that you don't have to invest your time and resources on creating a deep learning pipeline. Amazon Rekognition continues to improve the accuracy of its models by building upon the latest research and sourcing new training data. This allows you to focus on high-value application design and development.

### **Q: What are the most common use cases for Amazon Rekognition?**

The most common use-cases for Rekognition Image include:

- Searchable Image Library
- Face-Based User Verification
- Sentiment Analysis
- Facial Recognition
- Image Moderation

The most common use-cases for Rekognition Video include:

- Search Index for video archives
- Easy filtering of video for explicit and suggestive content

### **Q: How do I get started with Amazon Rekognition?**

If you are not already signed up for Amazon Rekognition, you can click the "Try Amazon Rekognition" button on the [Amazon Rekognition page](#) and complete the sign-up process. You must have an Amazon Web Services account; if you do not already have one, you will be prompted to create one during the sign-up process. Once you are signed up, try out Amazon Rekognition with your own images and videos using the [Amazon Rekognition Management Console](#) or download the [Amazon Rekognition SDKs](#) to start creating your own applications. Please refer to our [step-by-step Getting Started Guide](#) for more information.

### **Q: What image and video formats does Amazon Rekognition support?**

Amazon Rekognition Image currently supports the JPEG and PNG image formats. You can submit images either as an S3 object or as a byte array. Amazon Rekognition Video operations can analyze

videos stored in Amazon S3 buckets. The video must be encoded using the H.264 codec. The supported file formats are MPEG-4 and MOV. A codec is software or hardware that compresses data for faster delivery and decompresses received data into its original form. The H.264 codec is commonly used for the recording, compression and distribution of video content. A video file format may contain one or more codecs. If your MOV or MPEG-4 format video file does not work with Rekognition Video, check that the codec used to encode the video is H.264.

**Q: What file sizes can I use with Amazon Rekognition?**

Amazon Rekognition Image supports image file sizes up to 15MB when passed as an S3 object, and up to 5MB when submitted as an image byte array. Amazon Rekognition Video supports up to 10 GB files and up to 6 hour videos when passed through as an S3 file.

**Q: How does image resolution affect the quality of Rekognition Image API results ?**

Amazon Rekognition works across a wide range of image resolutions. For best results we recommend using VGA (640x480) resolution or higher. Going below QVGA (320x240) may increase the chances of missing faces, objects, or inappropriate content; although Amazon Rekognition accepts images that are at least 80 pixels in both dimensions.

**Q: How small can an object be for Amazon Rekognition Image to detect and analyze it?**

As a rule of thumb, please ensure that the smallest object or face present in the image is at least 5% of the size (in pixels) of the shorter image dimension. For example, if you are working with a 1600x900 image, the smallest face or object should be at least 45 pixels in either dimension.

**Q: How can I get Amazon Rekognition predictions reviewed by humans?**

Amazon Rekognition is directly integrated with [Amazon Augmented AI \(Amazon A2I\)](#) so you can easily route low confidence predictions from Amazon Rekognition Image to human reviewers. Using the Amazon Rekognition API for content moderation or the Amazon A2I console, you can specify the conditions under which Amazon A2I routes predictions to reviewers, which can be either a confidence threshold or a random sampling percentage. If you specify a confidence threshold, Amazon A2I routes only those predictions that fall below the threshold for human review. You can adjust these thresholds at any time to achieve the right balance between accuracy and cost-effectiveness. Alternatively, if you specify a sampling percentage, Amazon A2I routes a random sample of the predictions for human review. This can help you implement audits to monitor the prediction accuracy regularly. Amazon A2I also provides reviewers with a web interface consisting of all the instructions and tools they need to complete their review tasks. For more information

about implementing human review with Amazon Rekognition, see the [Amazon A2I webpage](#).

**Q: How does video resolution affect the quality of Rekognition Video API results?**

The system is trained to recognize faces larger than 32 pixels (on the shortest dimension), which translate into a minimum size for a face to be recognized that varies from approximately 1/7 of the screen smaller dimension at QVGA resolution to 1/30 at HD 1080p resolution. For example, at VGA resolution, users should expect lower performances for faces smaller than 1/10 of the screen smaller dimension.

**Q: What else can affect the quality of the Rekognition Video APIs ?**

Besides video resolution, heavy blur, fast moving persons, lighting conditions, pose may affect the quality of the APIs.

**Q: What is the preferred user video content that is suitable for Rekognition Video APIs?**

This API works best with consumer and professional videos taken with frontal field of view in normal color and lighting conditions. This API is not tested for black and white, IR or extreme lighting condition. Applications that are sensitive to false alarms are advised to discard outputs with confidence score below a selected (application-specific) confidence score.

**Q: In which AWS regions is Amazon Rekognition available?**

For a list of all regions where Amazon Rekognition is available, see the [AWS Region table](#).

## Label Detection

**Q: What is a label?**

A label is an object, scene, or concept found in an image based on its contents. For example, a photo of people on a tropical beach may contain labels such as 'Person', 'Water', 'Sand', 'Palm Tree', and 'Swimwear' (objects), 'Beach' (scene), and 'Outdoors' (concept).

**Q: What is a confidence score and how do I use it?**

A confidence score is a number between 0 and 100 that indicates the probability that a given prediction is correct. In the tropical beach example, if the object and scene detection process

returns a confidence score of 99 for the label 'Water' and 35 for the label 'Palm Tree', then it is more likely that the image contains water but not a palm tree.

Applications that are very sensitive to detection errors (false positives) should discard results associated with confidence scores below a certain threshold. The optimum threshold depends on the application. In many cases, you will get the best user experience by setting minimum confidence values higher than the default value.

### **Q: What is Object and Scene Detection?**

Object and Scene Detection refers to the process of analyzing an image or video to assign labels based on its visual content. Amazon Rekognition Image does this through the [DetectLabels API](#). This API lets you automatically identify thousands of objects, scenes, and concepts and returns a confidence score for each label. DetectLabels uses a default confidence threshold of 50. Object and Scene detection is ideal for customers who want to search and organize large image libraries, including consumer and lifestyle applications that depend on user-generated content and ad tech companies looking to improve their targeting algorithms.

### **Q: Can Amazon Rekognition detect object locations and return bounding boxes?**

Yes, Amazon Rekognition can detect the location of many common objects such as 'Person', 'Car', 'Gun', or 'Dog' in both images and videos. You get the coordinates of the bounding rectangle for each instance of the object found, as well as a confidence score. For more details on the API response structure for object bounding boxes, please refer to [the documentation](#).

### **Q: Does Amazon Rekognition provide information on the relationship between detected labels?**

Yes, for every label found, Amazon Rekognition returns its parent, alias, and category if they exist. Parents are returned in the "parents" field in hierarchical order. The first parent label is the immediate parent, while the following labels are parents of parents. For example, when a 'Car' is identified, Amazon Rekognition returns two parent labels 'Vehicle' (parent), and 'Transportation' (parent's parent). Aliases are labels with the same meaning as the primary labels and returned in the "aliases" field. For example, since 'Cell Phone' is an alias of 'Mobile Phone', Amazon Rekognition returns 'Cell Phone' in the "aliases" field of a 'Mobile Phone' label. Categories group labels based on common themes and returned in the "categories" field. For example, since 'Dog' is a label under the 'Animals and Pets' category, Amazon Rekognition returns 'Animal and Pets' in the "categories" field of a 'Dog' label. For more details on the full list of supported labels and their taxonomy, please visit [Amazon Rekognition Label Detection documentation](#).

### **Q: What types of labels does Amazon Rekognition support?**

Rekognition supports thousands of labels belonging to common categories including, but not limited to:

- People and Events: 'Wedding', 'Bride', 'Baby', 'Birthday Cake', 'Guitarist', etc.
- Food and Drink: 'Apple', 'Sandwich', 'Wine', 'Cake', 'Pizza', etc.
- Nature and Outdoors: 'Beach', 'Mountains', 'Lake', 'Sunset', 'Rainbow', etc.
- Animals and Pets: 'Dog', 'Cat', 'Horse', 'Tiger', 'Turtle', etc.
- Home and Garden: 'Bed', 'Table', 'Backyard', 'Chandelier', 'Bedroom', etc.
- Sports and Leisure: 'Golf', 'Basketball', 'Hockey', 'Tennis', 'Hiking', etc.
- Plants and Flowers: 'Rose', 'Tulip', 'Palm Tree', 'Forest', 'Bamboo', etc.
- Art and Entertainment: 'Sculpture', 'Painting', 'Guitar', 'Ballet', 'Mosaic', etc.
- Transportation and Vehicles: 'Airplane', 'Car', 'Bicycle', 'Motorcycle', 'Truck', etc.
- Electronics: 'Computer', 'Mobile Phone', 'Video Camera', 'TV', 'Headphones', etc.
- Landmarks: 'Brooklyn Bridge', 'Colosseum', 'Eiffel Tower', 'Machu Picchu', 'Taj Mahal', etc.

### **Q: How is Object and Scene Detection different for video analysis?**

Rekognition Video enables you to automatically identify thousands of objects - such as vehicles or pets - and activities - such as celebrating or dancing - and provides you with timestamps and a confidence score for each label. It also relies on motion and time context in the video to accurately identify complex activities, such as “blowing a candle” or “extinguishing fire”.

### **Q: I can't find the label I need. How do I request a new label?**

Please send us your label requests through the Amazon Rekognition Console by typing the label name in the input field of the 'Search all labels' section and click 'Request Rekognition to detect' the requested label. Amazon Rekognition continuously expands its catalog of labels based on customer feedback.

### **Q: What is Image Properties?**

Image Properties is a feature of Amazon Rekognition Image to detect dominant colors and image quality. Image Properties detects dominant colors of the entire image, image foreground, image background, and objects with localized bounding boxes. Image Properties also measure image quality through brightness, sharpness, and contrast scores. Image Properties can be called through



the DetectLabels API using IMAGE\_PROPERTIES as input parameter, with or without GENERAL\_LABEL input parameter for label detection. Visit the [Amazon Rekognition Label Detection documentation](#) to learn more.

**Q: How are dominant colors determined?**

Image Properties returns dominant colors in four formats: RGB, hexcode, CSS color, and simplified colors. Amazon Rekognition first identifies the dominant colors of by pixel percentage, and then maps these colors to the [140 CSS color palette](#), RGB, hex code, and 12 simplified colors (i.e., 'green', 'pink', 'black', 'red', 'yellow', 'cyan', 'brown', 'orange', 'white', 'purple', 'blue', 'grey'). By default, Image Properties returns ten (10) dominant colors unless customers specify the number of colors to return. The maximum number of dominant colors the API can return is 12.

**Q: How do I interpret the brightness, sharpness, and contrast scores?**

Image Properties provides a value ranging from 0 to 100 for each brightness, sharpness, and contrast score. For example, an underexposed image will return a low brightness score, while a brightly lit image will return a high brightness score.

**Q: How can you check if Amazon Rekognition has updated its models?**

Amazon Rekognition returns a LabelModelVersion parameter that lets you know whether the model has been updated. Object and Scene detection models are updated frequently based on customer feedback.

## Amazon Rekognition Custom Labels

**Q: Can I use Custom Labels for analyzing faces, customized text detection?**

No. Custom Labels is meant for finding objects and scenes in images. Custom Labels is not designed for analyzing faces, customized text detection. You should use other Rekognition APIs for these tasks. Please refer to the documentation for face analysis, Text detection.

**Q: Can I use Custom Labels for finding unsafe image content?**

Yes. Custom Labels is meant for finding objects and scenes in images. Custom Labels, when trained to detect your use case specific unsafe image content, can detect unsafe image content specific to your use case. Please also refer to the documentation for Moderation API to detect generic unsafe image content.

**Q: How many images are needed to train a custom model?**

The number of images required to train a custom model depends on the variability of the custom labels you want the model to predict and the quality of the training data. For example, a distinct logo overlaid on an image can be detected with 1-2 training images, while a more subtle logo required to be detected under many variations (scale, viewpoint, deformations) may need in the order of tens to hundreds of training examples with high quality annotations. If you already have a high number of labeled images, we recommend training a model with as many images as you have available. Please refer to the documentation for limits on maximum training dataset size.

Although hundreds of images may sometimes be required to train a custom model with high accuracy, with Custom Labels you can first train a model with tens of images per label, review your test results to understand where it does not work, and accordingly add new training images and train again to iteratively improve your model.

**Q: How many inference compute resources should I provision for my custom model?**

The number of parallel inference compute resources needed depends on how many images you need to process at a given point in time. The throughput of a single resource will depend factors including the size of the images, the complexity of those images (how many detected objects are visible), and the complexity of your custom model. We recommend that you monitor the frequency at which you need provision your custom model, and the number of images that need to be processed at a single time, in order to schedule provisioning of your custom model most efficiently. If you expect to process images periodically (e.g. once a day or week, or scheduled times during the day), you should Start provisioning your custom model at a scheduled time, process all your images, and then Stop provisioning. If you don't stop provisioning, you will be charged even if no images are processed.

**Q: My training has failed. Will I be charged?**

No. You will not be charged for the compute resources if your training fails.

## Content Moderation

**Q: What is Content Moderation?**

Amazon Rekognition's Content Moderation API uses deep learning to detect explicit or suggestive adult content, violent content, weapons, visually disturbing content, drugs, alcohol, tobacco, hate

symbols, gambling, and rude gestures in image and videos. Beyond flagging an image or video based on presence of inappropriate or offensive content, Amazon Rekognition also returns a hierarchical list of labels with confidence scores. These labels indicate specific sub-categories of the type of content detected, thus providing more granular control to developers to filter and manage large volumes of user generated content (UGC). This API can be used in moderation workflows for applications such as social and dating sites, photo sharing platforms, blogs and forums, apps for children, e-commerce site, entertainment and online advertising services.

**Q: What types of inappropriate, offensive, and unwanted content does Amazon Rekognition detect?**

You can find a full list of content categories detected by Amazon Rekognition [here](#).

Amazon Rekognition returns a hierarchy of labels, as well as a confidence score for each detected label. For instance, given an inappropriate image, Rekognition may return “Explicit Nudity” with a confidence score as a top-level label. Developers can use this metadata to flag content at a high level, for example, when all types of explicit adult content is to be flagged. In the same response, Rekognition also returns second level of granularity by providing additional context like “Graphic Male Nudity” with its own confidence score. Developers can use this information to build more complex filtering logic to serve different geographies and demographics.

Please note that the Content Moderation API is not an authority on, or in any way purports to be an exhaustive filter of inappropriate and offensive content. Furthermore, this API does not detect whether an image includes illegal content (such as child pornography) or unnatural adult content.

If you require other types of inappropriate content to be detected in images, please reach out to us using the feedback process outlined later in this section.

**Q: How can I know which model version I am currently using?**

Amazon Rekognition makes regular improvement to its models. To keep track of model version, you can use the 'ModerationModelVersion' field in the API response.

**Q: How can I ensure that Amazon Rekognition meets accuracy goals for my image or video moderation use case?**

Amazon Rekognition’s Content Moderation models have been and tuned and tested extensively, but we recommend that you measure the accuracy on your own data sets to gauge performance.

You can use the 'MinConfidence' parameter in your API requests to balance detection of content (recall) vs the accuracy of detection (precision). If you reduce 'MinConfidence', you are likely to detect most of the inappropriate content, but are also likely to pick up content that is not actually inappropriate. If you increase 'MinConfidence' you are likely to ensure that all your detected content is truly inappropriate but some content may not be tagged.

### **Q: How can I give feedback to Rekognition to improve its Content Moderation APIs?**

Please send us your requests through [AWS Customer Support](#). Amazon Rekognition continuously expands the types of inappropriate content detected based on customer feedback. Please note that illegal content (such as child pornography) will not be accepted through this process.

## **Facial Analysis**

### **Q: What is Facial Analysis?**

Facial analysis is the process of detecting a face within an image and extracting relevant face attributes from it. Amazon Rekognition Image takes returns the bounding box for each face detected in an image along with attributes such as gender, presence of sunglasses, and face landmark points. Rekognition Video will return the faces detected in a video with timestamps and, for each detected face, the position and a bounding box along with face landmark points.

### **Q: What face attributes can I get from Amazon Rekognition?**

Amazon Rekognition returns the following facial attributes for each face detected, along with a bounding box and confidence score for each attribute:

- Gender
- Smile
- Emotions
- Eyeglasses
- Sunglasses
- Eyes open
- Mouth open
- Mustache

- Beard
- Pose
- Quality
- Face landmarks

**Q: What is face pose?**

Face pose refers to the rotation of a detected face on the pitch, roll, and yaw axes. Each of these parameters is returned as an angle between -180 and +180 degrees. Face pose can be used to find the orientation of the face bounding polygon (as opposed to a rectangular bounding box), to measure deformation, to track faces accurately, and more.

**Q: What is face quality?**

Face quality describes the quality of the detected face image using two parameters: sharpness and brightness. Both parameters are returned as values between 0 and 1. You can apply a threshold to these parameters to filter well-lit and sharp faces. This is useful for applications that benefit from high-quality face images, such as face comparison and face recognition.

**Q: What are face landmarks?**

Face landmarks are a set of salient points, usually located on the corners, tips or mid points of key facial components such as the eyes, nose, and mouth. Amazon Rekognition [DetectFaces API](#) returns a set of face landmarks that can be used to crop faces, morph one face into another, overlay custom masks to create custom filters, and more.

**Q: How many faces can I detect in an image?**

You can detect up to 100 faces in an image using Amazon Rekognition.

**Q: How is Facial Analysis different for video analysis?**

With Rekognition Video, you can locate faces across a video and analyze face attributes, such as whether the face is smiling, eyes are open, or showing emotions. Rekognition Video will return the detected faces with timestamps and, for each detected face, the position and a bounding box along with landmark points such as left eye, right eye, nose, left corner of the mouth, and right corner of the mouth. This position and time information can be used to easily track user sentiment over time and deliver additional functionality such as automatic face frames, highlights, or crops.

## **Q: In addition to Video resolution, what else can affect the quality of the Rekognition Video APIs?**

Besides video resolution, the quality and representative faces, part of the face collections to search, has major impact. Using multiple face instances per person with variations like beard, glasses, poses (profile and frontal) will significantly improve the performance. Typically very fast moving people and blurred videos may experience lower quality.

## **Face Comparison**

### **Q: What is Face Comparison?**

Face Comparison is the process of comparing one face to one or more faces to measure similarity. Using the [CompareFaces API](#), Amazon Rekognition Image lets you measure the likelihood that faces in two images are of the same person. The API compares a face in the source input image with each face detected in the target input image and returns a similarity score for each comparison. You also get a bounding box and confidence score for each face detected. You can use face comparison to verify a person's identity against their personnel photo on file in near real-time.

### **Q: Can I use a source image with more than one face?**

Yes. If the source image contains multiple faces, [CompareFaces](#) detects the largest face and uses it to compare with each face detected in the target image.

### **Q: How many faces can I compare against?**

You can compare one face in the source image with up to 15 detected faces in the target image.

## **Facial Recognition**

### **Q: What is Facial Recognition?**

Facial recognition is the process of identifying or verifying a person's identity by searching for their face in a collection of faces. Using facial recognition, you can easily build applications such as multi-factor authentication for bank payments, automated building entry for employees, and more.

**Q: What is a face collection and how do I create one?**

A face collection is a searchable index of face feature vectors, owned and managed by you. Using the [CreateCollection](#) API, you can easily create a collection in a supported AWS region and get back an Amazon Resource Name (ARN). Each face collection has a unique CollectionId associated with it.

**Q: How do I add faces to or delete faces from a face collection?**

To add a face to an existing face collection, use the [IndexFaces](#) API. This API accepts an image in the form of an S3 object or image byte array and adds a vector representation of the faces detected to the face collection. IndexFaces also returns a unique Faceld and face bounding box for each of the faces added.

To delete a face from an existing face collection, use the [DeleteFaces](#) API. This API operates on the face collection supplied (using a CollectionId) and removes the entries corresponding to the list of Facelds. For more information on adding and deleting faces, please refer to our [Managing Collections](#) example.

**Q: How do I search for a face within a face collection?**

Once you have created an indexed collection of faces, you can search for a face within it using either an image ([SearchFaceByImage](#)) or a Faceld ([SearchFaces](#)). These APIs take in an input face and return a set of faces that match, ordered by similarity score with the highest similarity first. For more details, please refer to our [Searching Faces](#) example.

**Q: How is Facial Recognition different for video analysis?**

Rekognition Video allows you to perform real time face searches against collections with tens of millions of faces. First, you create a face collection, where you can store faces, which are vector representations of facial features. Rekognition then searches the face collection for visually similar faces throughout your video. Rekognition will return a confidence score for each of the faces in your video, so you can display likely matches in your application.

**Q: In addition to Video resolution what else can affect the quality of the Video APIs ?**

Besides video resolution, the quality and representative faces part of the face collections to search has major impact. Using multiple face instances per person with variations like beard, glasses, poses (profile and frontal) will significantly improve the performance. Typically very fast moving people may experience low recall. In addition, blurred videos may also experience lower quality.

# Celebrity Recognition

## **Q: What is Celebrity Recognition?**

Amazon Rekognition's Celebrity Recognition is a deep learning based easy-to-use API for detection and recognition of individuals who are famous, noteworthy, or prominent in their field. The RecognizeCelebrities API has been built to operate at scale and recognize celebrities across a number of categories, such as politics, sports, business, entertainment, and media. Our Celebrity Recognition feature is ideal for customers who need to index and search their digital image libraries for celebrities based on their particular interest.

## **Q: Who can be identified by the Celebrity Recognition API?**

Amazon Rekognition can only identify celebrities that the deep learning models have been trained to recognize. Please note that the RecognizeCelebrities API is not an authority on, and in no way purports to be, an exhaustive list of celebrities. The feature has been designed to include as many celebrities as possible, based on the needs and feedback of our customers. We are constantly adding new names, but the fact that Celebrity Recognition does not recognize individuals that may be deemed prominent by any other groups or by our customers is not a reflection of our opinion of their celebrity status. If you would like to see additional celebrities identified by Celebrity Recognition, please submit feedback.

## **Q: Can a celebrity identified through the Amazon Rekognition API request to be removed from the feature?**

Yes. If a celebrity wishes to be removed from the feature, he or she can send an email to AWS Customer Support and we will process the removal request.

## **Q: What sources are supported to provide additional information about a Celebrity ?**

The API supports an optional list of sources to provide additional information about the celebrity as a part of the API response. We currently provide the IMDB URL, when it is available. We may add other sources at a later date.

## **Q: How is Celebrity Recognition different for video analysis?**

With Rekognition Video, you can detect and recognize when and where well known persons appear in a video. The time-coded output includes the name and unique id of the celebrity, bounding box coordinates, confidence score, and URLs pointing to related content for the celebrity, for example,



the celebrity's IMDB link. The celebrity is also detected even if sometimes the face becomes occluded in the video. This feature allows you to index and search digital video libraries for use cases related to your specific marketing and media needs.

**Q: In addition to Video resolution, what else can affect the quality of the Rekognition Video APIs?**

Very fast moving celebrities and blurred videos can affect the quality of the Rekognition Video APIs. In addition, heavy makeup and camouflage common for actors/actresses, can also affect the quality.

## Text Detection

**Q: What is Text Detection?**

Text detection is a capability of Amazon Rekognition that allows you to detect and recognize text within an image or a video, such as street names, captions, product names, overlaid graphics, video subtitles, and vehicular license plates. Text detection is specifically built to work with real-world images and videos, rather than document images. Amazon Rekognition's DetectText API takes in an image and returns the text label and a bounding box for each detected string of characters, along with a confidence score. For example, in image sharing and social media applications, you can enable visual search based on an index of images that contain the same text labels. In security applications, you can identify vehicles based on license plate numbers from images taken by traffic cams. Similarly, for videos, using the StartTextDetection and GetTextDetection APIs, you can detect text and get confidence scores and timestamps for each detection. In media and entertainment applications, you can create text metadata to support search for relevant content, such as news, sport scores, commercials, and captions. You can also review the detected text for policy or compliance violations e.g. an email address or phone number that has been overlaid by spammers.

**Q: What type of text does Amazon Rekognition Text Detection support?**

Text Detection is specifically built to work with real-world images and videos rather than document images. It supports text in most Latin scripts and numbers embedded in a large variety of layouts, fonts and styles, and overlaid on background objects at various orientation as banners and posters. Text detection recognizes up to 50 sequences of characters per the image or video frame and lists them as words and lines. Text detection supports text rotated by up to -90 to +90 degrees from the horizontal axis.

**Q: Can I limit text detection to specific regions in an image or video frame?**

Yes, you can use text detection filtering options to specify up to 10 Regions of Interest (ROIs) in the API request. Amazon Rekognition will only return text that falls within these regions.

**Q: Can I filter text detections by word confidence or bounding box size?**

Yes, in the API request you can use the text detection filtering options to specify thresholds for minimum confidence scores or minimum bounding box dimensions.

**Q: How can I give feedback to Rekognition to improve its text recognition?**

Please send us your requests through [AWS Customer Support](#). Amazon Rekognition continuously expands the types of text content recognized based on customer feedback.

## PPE Detection

**Q: What personal protective equipment (PPE) can Amazon Rekognition detect?**

Amazon Rekognition “DetectProtectiveEquipment” can detect common types of face covers, hand covers, and head covers. To learn more, please refer to [feature documentation](#). You can also use Amazon Rekognition Custom Labels to detect PPE such as high-visibility vests, safety goggles, and other PPE unique to your business. To learn about how you can use Amazon Rekognition Custom Labels for custom PPE detection, visit this [github repo](#).

**Q: Can Amazon Rekognition detect protective equipment locations and return bounding boxes?**

Yes, Amazon Rekognition “DetectProtectiveEquipment” API can detect the location of protective equipment such as face covers, hand covers, and head covers on persons in images. You get the coordinates of the bounding box rectangle for each item of protective equipment detected, as well as a confidence score. For more details on the API response, please refer to [documentation](#).

**Q: Can the service detect if the mask is worn properly?**

Amazon Rekognition “DetectProtectiveEquipment” API output provides “CoversBodyPart” value (true/false) and confidence value for the Boolean value for each detected item of protective equipment. This provides information on whether the protective equipment is on the corresponding body part of the person. The prediction about the presence of the protective equipment on the corresponding body part helps filter out cases where the PPE is in the image but

not actually on the person. It does not, however, indicate or imply that the person is adequately protected by the protective equipment or that the protective equipment itself is properly worn.

**Q: Can Amazon Rekognition PPE detection identify detected persons?**

No, Amazon Rekognition PPE detection does not perform facial recognition or facial comparison and cannot identify the detected persons.

**Q: Where can I find more information about API limits and latency?**

Please refer to [Amazon Rekognition PPE detection documentation](#) to get the latest details on API limits and latency.

**Q: How do I send images from my workplace cameras to Amazon Rekognition?**

You have multiple options to sample images from your workplace cameras. Please refer to the [Amazon Rekognition PPE detection blog](#) to learn more.

**Q: How is PPE detection priced?**

Amazon Rekognition PPE detection is priced similarly to other Amazon Rekognition Image APIs on a per image basis. To learn more, visit the [Amazon Rekognition pricing page](#).

## Amazon Rekognition Streaming Video Events

**Q: What are Amazon Rekognition Streaming Video Events?**

Amazon Rekognition Streaming Video Events uses machine learning to detect objects from connected camera to provide actionable alerts in real time. Amazon Rekognition Streaming Video events work with your new and existing Kinesis Video Streams to process video streams (up to 120 seconds per motion event) and notify you as soon a desired object of interest is detected. You can use these notifications to

- Send Smart Alerts to your end users such as “a package was detected at the front door.”
- Provide home automation capabilities such as “turning on the garage light when a person is detected.”
- Integrate with smart assistants such as Echo devices to provide Alexa announcements when an object is detected.

- Provide Smart Search capabilities such as search for all video clips where a package was detected.

**Q: How does Amazon Rekognition Streaming Video Events work?**

You can use your new or existing Kinesis Video Streams to get started with Amazon Rekognition Streaming Video events. When configuring your stream processor settings for Amazon Rekognition you can choose the desired labels (person, pet, or package) that you want to detect, the duration for the video (up to 120 seconds per motion event) that Rekognition should process for each event, and/or the region of interest on the frame that you want to process through Rekognition. Rekognition Streaming Video Event APIs process video only when you send a notification to Rekognition to start processing the video stream.

When motion is detected on a connected camera, you send a notification to Rekognition to start processing the video stream. Rekognition processes the corresponding Kinesis Video Stream, post motion detection, to look for the desired objects specified by you. As soon as a desired object is detected, Amazon Rekognition will send you a notification. This notification includes the object detected, bounding box, zoomed in image of the object, and the time stamp.

**Q: What labels can Amazon Rekognition Streaming Video Events support?**

Amazon Rekognition Streaming Video Events can support people, pets, and packages.

**Q: What pets and package types can Amazon Rekognition Streaming Video APIs detect?**

Amazon Rekognition Streaming Video Event APIs support dogs and cats for pet detection. The API can detect medium and large cardboard boxes with high accuracy. The API also detects smaller boxes, bubble mailer envelopes, and folders but may miss some of these objects occasionally.

**Q: Will I be charged separately for each label detected? Can I choose which labels to opt into?**

No, you will not be charged separately for each label. You will be charged for the duration of streaming video processed by Rekognition. You can either opt into specific labels (pet, package) or choose to opt in to all three labels (people, pet, package) while configuring your stream processing settings.

**Q: Do I need to stream video continuously to Amazon Rekognition?**

No, you do not have to stream video continuously to Amazon Rekognition.

**Q: Should I create new Kinesis Video Streams (KVS) to use Streaming Video Events?**

Amazon Rekognition Streaming Video Events works with both new and existing Kinesis Video Streams. Simply integrate the relevant KVS streams with Amazon Rekognition Streaming Video Events API to get started with video analysis on KVS streams.

**Q: When will Amazon Rekognition send me a notification?**

Amazon Rekognition starts processing the video stream post motion detection. You can configure the duration for processing this video stream (up to 120 seconds per event). As soon as Amazon Rekognition detects the object of interest in the video stream, Rekognition will send you a notification. This notification includes the type of object detected, the bounding box, a zoomed in image of the object detected, and a time stamp.

**Q: What resolution and fps is support for label detection?**

In order to keep costs and latency low Amazon Rekognition Streaming Video Events support 1080p or lower resolution video streams. Rekognition processes the video stream at 5 fps.

**Q: What codecs and file format is supported for streaming video?**

Amazon Rekognition Video supports H.264 files in MPEG-4 (.mp4) or MOV format.

**Q: What is the maximum duration of the video processed per event?**

You can process up to 120 seconds of video per event.

**Q: Can I choose a particular area of the frame to be processed for my video stream?**

Yes, as a part of configuring your StreamProcessor you can choose the region of interest that you want to process on your frame. Amazon Rekognition will only process that particular area of the frame.

**Q: How many concurrent video streams can I process with Amazon Rekognition?**

Amazon Rekognition Streaming Video Events can support 600 concurrent sessions per AWS customer. Please reach out to your account manager if you need to increase this limit.

## **Amazon Rekognition Stored Video Analysis**

**Q: What types of entities can Amazon Rekognition Video detect?**

Amazon Rekognition Video can detect objects, scenes, landmarks, faces, celebrities, text, and inappropriate content in videos. You can also search for faces appearing in a video using your own repository or collection of face images.

**Q: What types file formats and codecs does Amazon Rekognition Video support?**

Amazon Rekognition Video supports H.264 files in MPEG-4 (.mp4) or MOV format. If your video files use a different codec, you can transcode them into H.264 using [AWS Elemental MediaConvert](#).

**Q: How do Amazon Rekognition Video asynchronous APIs work?**

Amazon Rekognition Video can process videos stored in an Amazon S3 bucket. You can use an asynchronous set of operations: you start video analysis by calling a Start operation such as `StartLabelDetection` for detecting objects and scenes. The completion status of the request is published to an Amazon Simple Notification Service (SNS) topic. To get the completion status from the Amazon SNS topic, you can use an Amazon Simple Queue Service (SQS) queue or an AWS Lambda function. Once you have the completion status, you call a Get operation such as `GetLabelDetection` to get the results of the request. For a list of available Amazon Rekognition Video APIs, please see [this page](#).

**Q: How can I find the timeline for each detection in a video?**

Amazon Rekognition Video returns label results by timestamps or video segments. You can choose how you want to organize these results by using the `AggregateBy` input parameter in `GetLabelDetection` API.

- When label results are organized by timestamps, each label will be returned every time Amazon Rekognition Video detects that label in the video timeline. For example, if 'Dog' is detected at 2000ms and 4000ms, Amazon Rekognition Video will return 2 label entries for 'Dog', one with 2000ms and another with 4000ms.
- When label results are organized by video segments, Amazon Rekognition Video returns the video segment for when a label is detected across multiple consecutive frames. A video segment is defined by a start timestamp, an end timestamp, and a duration. For example, if 'Dog' is detected in 2 consecutive frames at 2000ms and 4000ms, Amazon Rekognition Video will return 1 label entry for 'Dog' with start timestamp 2000ms, end timestamp 4000ms, and duration 2000ms.

To learn more about timestamps and segments and see a sample API response, visit [Detecting labels in a video](#).

**Q: How many concurrent video analysis jobs can I run with Amazon Rekognition Video?**

You can process up to 20 concurrent jobs with Amazon Rekognition Video. For more details on limits, please see our [limits page](#).

**Q: What video resolution should I use?**

Amazon Rekognition Video automatically handles a wide range of video resolutions and video quality. We recommend using a 720p (1280×720 pixels) to 1080p (1920×1080 pixels), or their equivalent resolutions in other aspect ratios for optimum results. Very low resolution (such as QVGA or 240p) and very low-quality videos may adversely impact results quality.

## Q: What is People Pathing?

With Rekognition Video, you can find the path of each person across the video timeline. Rekognition Video detects persons even when the camera is in motion and, for each person, returns a bounding box and the face, along with face attributes and timestamps. For retail applications, this allows to generate customer insights anonymously, such as how customers move across aisles in a shopping mall or how long they are waiting in checkout lines.

## Media Analysis using Amazon Rekognition Video

### Q: What types of media analysis segments can Amazon Rekognition Video detect?

Amazon Rekognition Video can detect the following types of segments or entities for media analysis:

- **Black frames:** Videos often contain a short duration of empty black frames with no audio that are used as cues to insert advertisements, or to demarcate the end of a program segment such as a scene or the opening credits. With Amazon Rekognition Video, you can detect such black frame sequences to automate ad insertion, package content for VOD, and demarcate various program segments or scenes. Black frames with audio (such as fade outs or voiceovers) are considered as content and not returned.
- **Credits:** Amazon Rekognition Video helps you automatically identify the exact frames where the opening and closing credits start and end for a movie or TV show. With this information, you can generate 'binge markers' or interactive viewer prompts such as 'Next Episode' or 'Skip Intro' in VOD applications. Amazon Rekognition Video is trained to handle a wide variety of opening and end credit styles ranging from simple rolling credits to more challenging credits alongside content, credits on scenes, or stylized credits in anime content.
- **Shots:** A shot is a series of interrelated consecutive pictures taken contiguously by a single camera and representing a continuous action in time and space. With Amazon Rekognition Video, you can detect the start, end, and duration of each shot, as well as a count all the shots in a piece of content. Shot metadata can be used for applications such as creating promotional videos using selected shots, generating a set of preview thumbnails that avoid transitional content between shots, and inserting ads in spots that don't disrupt viewer experience, such as the middle of a shot when someone is speaking.
- **Color Bars:** Amazon Rekognition Video allows you to detect sections of video that display SMPTE or EBU color bars, which are a set of colors displayed in specific patterns to ensure color is calibrated correctly on broadcast monitors, programs, and on cameras. For more information



about SMPTE color bars, see [SMPTE color bar](#). This metadata is useful to prepare content for VOD applications by removing color bar segments from the content, or to detect issues such as loss of broadcast signals in a recording, when color bars are shown continuously as a default signal instead of content.

- **Slates:** Slates are sections, typically at the beginning of a video, that contain text metadata about the episode, studio, video format, audio channels, and more. Amazon Rekognition can identify the start and end such slates, making it easy for operators to use the text metadata or to simply remove the slate when preparing content for final viewing.
- **Studio logos:** Studio logos are sequences that show the logos or emblems of the production studio involved in making the show. Amazon Rekognition can identify such sequences, making it easy for operators to review them for identifying studios.
- **Content:** Content refers to the portions of the TV show or movie that contain the program or related elements. Black frames, credits, color bars, slates, and studio logos are not considered to be content. Amazon Rekognition Video enables you to detect the start and end of each content segment in the video, which enables multiple uses such as finding the program run time or finding certain segments that serve specific purposes. For example, a quick recap of the previous episode at the beginning of the video is a type of content. Similarly, bonus post-credit content can appear after the credits have finished. And, some videos may have 'textless' content at the end of the video, which are a set of all program content that contains overlaid text, but with that text removed to enable internationalization in another language. Once all the content segments are detected with Amazon Rekognition Video, you can apply specific domain knowledge such as 'my videos always start with a recap' to further categorize each segment or to send them for human review.

Amazon Rekognition Video provides the start, end, duration, and timecodes for each detected entity, and provides timestamp (milliseconds), SMPTE format code, and frame number options for each.

### Q: How do I get started with media analysis using Amazon Rekognition Video?

Media analysis features are available through the Amazon Rekognition Video [segment detection API](#). This is an [asynchronous API](#) composed of two operations: [StartSegmentDetection](#) to start the analysis, and [GetSegmentDetection](#) to get the analysis results. To get started, please refer to the [documentation](#).

If you want to visualize the results of media analysis or even try out other Amazon AI services like Amazon Transcribe with your own videos, please use the [Media Insights application](#) – a serverless framework and demo application to easily generate insights and develop applications for your



video, audio, text, and image resources, using AWS Machine Learning and Media services. You can easily spin up your own demo application using the supplied AWS CloudFormation template, to try out your own videos and visualize analysis results.

**Q: What is a frame accurate timecode?**

Frame accurate timecodes provide the exact frame number for a relevant segment of video or entity. Media companies commonly process timecodes using the SMPTE (Society of Motion Picture and Television Engineers) format hours:minutes:seconds:frame number, for example, 00:24:53:22.

**Q: Is Amazon Rekognition Video segment detection frame accurate?**

Yes, the Amazon Rekognition Video segment detection API provides frame accurate SMPTE timecodes, as well as millisecond timestamps for the start and end of each detection.

**Q: What types of frame rate formats can Amazon Rekognition Video segment detection handle?**

Amazon Rekognition Video segment detection automatically handles integer, fractional and drop frame standards for frame rates between 15 and 60fps. For example, common frame rates such as 23.976 fps, 25fps, 29.97 fps and 30fps are supported by segment detection. Frame rate information is utilized to provide frame accurate timecodes in each case.

**Q: What filtering options can I apply?**

You can specify the minimum confidence for each segment type while making the API request. For example, you can filter out any segment below 70% confidence score. For black frame detection, you can also control the maximum pixel luminance that you consider to be a black pixel, for example, a value of 40 for a color range of 0 to 255. Further, you can also control what percentage of pixels in a frame need to meet this black pixel luminance criteria for the frame to be classified as a black frame, for example, 99%. These filters allow you to account for varied video quality and formats when detecting black frames. For example, videos reclaimed from tape archives might be noisy and have a different black level compared to a modern digital video. For more details, please refer to [this page](#).

## Billing

**Q: How does Amazon Rekognition count the number of images processed?**

For APIs that accept images as inputs, Amazon Rekognition counts the actual number of images analyzed as the number of images processed. DetectLabels, DetectModerationLabels, DetectFaces, IndexFaces, RecognizeCelebrities, and SearchFaceByImage, and Image Properties belong to this category. For the CompareFaces API, where two images are passed as input, only the source image is counted as a unit of images processed.

For API calls that don't require an image as an input parameter, Amazon Rekognition counts each API call as one image processed. SearchFaces belongs to this category.

The remaining Amazon Rekognition APIs - ListFaces, DeleteFaces, CreateCollection, DeleteCollection, and ListCollections - do not count towards images processed.

**Q: How does Amazon Rekognition count the number of minutes of videos processed?**

For archived videos, Amazon Rekognition counts the minutes of video that is successfully processed by the API and meters them for billing. For Live stream videos you get charged in chunks of every five seconds of video that we successfully process.

**Q: Which APIs does Amazon Rekognition charge for?**

Amazon Rekognition Image charges for the following APIs: DetectLabels, DetectModerationLabels, DetectText, DetectFaces, IndexFaces, RecognizeCelebrities, SearchFaceByImage, CompareFaces, SearchFaces, and Image Properties. Amazon Rekognition Video charges are based on duration of video in minutes, successfully processed by StartLabelDetection, StartFaceDetection, StartFaceDetection, StartTextDetection, StartContentModeration, StartPersonTracking, StartCelebrityRecognition, StartFaceSearch and StartStreamProcessor APIs.

**Q: How much does Amazon Rekognition cost?**

Please see the [Amazon Rekognition Pricing Page](#) for current pricing information.

**Q: Will I be charged for the feature vectors I store in my face collections?**

Yes. Amazon Rekognition charges \$0.01 per 1,000 face vectors per month. For details, please see the [pricing page](#).

**Q: Does Amazon Rekognition participate in the AWS Free Tier?**

Yes. As part of the [AWS Free Usage Tier](#), you can get started with Amazon Rekognition for free. Upon sign-up, new Amazon Rekognition customers can analyze up to 5,000 images for free each

month for the first 12 months. You can use all Amazon Rekognition APIs, except for Image Properties, with this free tier, and also store up to 1,000 faces without any charge. In addition, Amazon Rekognition Video customers can analyze 1,000 minutes of Video free, per month, for the first year.

**Q: Do your prices include taxes?**

For details on taxes, please see [Amazon Web Services Tax Help](#).

## AWS Integration

**Q: Does Amazon Rekognition Video work with images stored on Amazon S3?**

Yes. You can start analyzing images stored in Amazon S3 by simply pointing the Amazon Rekognition API to your S3 bucket. You don't need to move your data. For more details of how to use S3 objects with Amazon Rekognition API calls, please see our [Detect Labels exercise](#).

**Q: Can I use Amazon Rekognition with images stored in an Amazon S3 bucket in another region?**

No. Please ensure that the Amazon S3 bucket you want to use is in the same region as your Amazon Rekognition API endpoint.

**Q: How do I process multiple image files in a batch using Amazon Rekognition?**

You can process your Amazon S3 images in bulk using the steps described in our [Amazon Rekognition Batch Processing example on GitHub](#).

**Q: How can I use AWS Lambda with Amazon Rekognition?**

Amazon Rekognition provides seamless access to AWS Lambda and allows you bring trigger-based image analysis to your AWS data stores such as Amazon S3 and Amazon DynamoDB. To use Amazon Rekognition with AWS Lambda, please follow the steps outlined [here](#) and select the Amazon Rekognition blueprint.

**Q: Does Amazon Rekognition work with AWS CloudTrail?**

Yes. Amazon Rekognition supports logging the following actions as events in CloudTrail log files: CreateCollection, DeleteCollection, CreateStreamProcessor, DeleteStreamProcessor,

DescribeStreamProcessor, ListStreamProcessors, and ListCollections. For more details on the Amazon Rekognition API calls that are integrated with AWS CloudTrail, see [Logging Amazon Rekognition API Calls with AWS CloudTrail](#).

## Data Privacy

### **Q: Are image and video inputs processed by Amazon Rekognition stored, and how are they used by AWS?**

Amazon Rekognition may store and use image and video inputs processed by the service solely to provide and maintain the service and, unless you opt out as provided below, to improve and develop the quality of Amazon Rekognition and other Amazon machine-learning/artificial-intelligence technologies. Use of your content is important for continuous improvement of your Amazon Rekognition customer experience, including the development and training of related technologies. We do not use any personally identifiable information that may be contained in your content to target products, services or marketing to you or your end users. Your trust, privacy, and the security of your content are our highest priority and we implement appropriate and sophisticated technical and physical controls, including encryption at rest and in transit, designed to prevent unauthorized access to, or disclosure of, your content and ensure that our use complies with our commitments to you. Please see <https://aws.amazon.com/compliance/data-privacy-faq/> for more information. You may opt out of having your image and video inputs used to improve or develop the quality of Amazon Rekognition and other Amazon machine-learning/artificial-intelligence technologies by using an AWS Organizations opt-out policy. For information about how to opt out, see [Managing AI services opt-out policy](#).

### **Q: Can I delete image and video inputs stored by Amazon Rekognition?**

Yes. You can request deletion of image and video inputs associated with your account by contacting [AWS Support](#). Deleting image and video inputs may degrade your Amazon Rekognition experience.

### **Q: Who has access to my content that is processed and stored by Amazon Rekognition?**

Only authorized employees will have access to your content that is processed by Amazon Rekognition. Your trust, privacy, and the security of your content are our highest priority and we implement appropriate and sophisticated technical and physical controls, including encryption at rest and in transit, designed to prevent unauthorized access to, or disclosure of, your content and

ensure that our use complies with our commitments to you. Please see <https://aws.amazon.com/compliance/data-privacy-faq/> for more information.

**Q: Do I still own my content that is processed and stored by Amazon Rekognition?**

You always retain ownership of your content and we will only use your content with your consent.

**Q: Is the content processed by Amazon Rekognition moved outside the AWS region where I am using Amazon Rekognition?**

Any content processed by Amazon Rekognition is encrypted and stored at rest in the AWS region where you are using Amazon Rekognition. Unless you opt out as provided below, some portion of content processed by Amazon Rekognition may be stored in another AWS region solely in connection with the continuous improvement and development of your Amazon Rekognition customer experience and other Amazon machine-learning/artificial-intelligence technologies. You can request deletion of image and video inputs associated with your account by contacting AWS Support. Your trust, privacy, and the security of your content are our highest priority and we implement appropriate and sophisticated technical and physical controls, including encryption at rest and in transit, designed to prevent unauthorized access to, or disclosure of, your content and ensure that our use complies with our commitments to you. Please see <https://aws.amazon.com/compliance/data-privacy-faq/> for more information. Your content will not be stored in another AWS region if you opt out of having your content used to improve and develop the quality of Amazon Rekognition and other Amazon machine-learning/artificial-intelligence technologies. For information about how to opt out, see [Managing AI services opt-out policy](#).

**Q: Can I use Amazon Rekognition in connection with websites, programs or other applications that are directed or targeted to children under age 13 and subject to the Children's Online Privacy Protection Act (COPPA)?**

Yes, subject to your compliance with the Amazon Rekognition Service Terms, including your obligation to provide any required notices and obtain any required verifiable parental consent under COPPA, you may use Amazon Rekognition in connection with websites, programs, or other applications that are directed or targeted, in whole or in part, to children under age 13.

**Q: How do I determine whether my website, program, or application is subject to COPPA?**

For information about the requirements of COPPA and guidance for determining whether your website, program, or other application is subject to COPPA, please refer directly to the resources provided and maintained by the [United States Federal Trade Commission](#). This site also contains

information regarding how to determine whether a service is directed or targeted, in whole or in part, to children under age 13.

### **Q: Is Amazon Rekognition a HIPAA Eligible Service?**

Amazon Rekognition is a HIPAA Eligible Service covered under the AWS Business Associate Addendum (AWS BAA). If you have an AWS BAA in place, Amazon Rekognition will use, disclose, and maintain your Protected Health Information (PHI) only as permitted by the terms of your AWS BAA.

## **Access Control**

### **Q: How do I control user access for Amazon Rekognition?**

Amazon Rekognition is integrated with [AWS Identity and Access Management \(IAM\)](#). AWS IAM policies can be used to ensure that only authorized users have access to Amazon Rekognition APIs. For more details, please see the [Amazon Rekognition Authentication and Access Control page](#).

## **Report Abuse**

### **Q: How can I report potential Amazon Rekognition abuse?**

If you suspect that Amazon Rekognition is being used in manner that is abusive or illegal, or infringes on your rights or the rights of other people, [please report this use and AWS will investigate the issue](#).

## **Learn more about Amazon Rekognition pricing**

[Visit the pricing page](#)

Ready to build?

[Get started with AWS Rekognition](#)

Have more questions?

[Contact us](#)**Get Started with AWS**

Learn how to start using AWS in minutes

**AWS Free Tier**

Gain free, hands-on experience with AWS for 12 months

[Sign In to the Console](#)

## Learn About AWS

[What Is AWS?](#)[What Is Cloud Computing?](#)[AWS Diversity, Equity & Inclusion](#)[What Is DevOps?](#)[What Is a Container?](#)[What Is a Data Lake?](#)[AWS Cloud Security](#)[What's New](#)[Blogs](#)[Press Releases](#)

## Help

[Contact Us](#)

## Resources for AWS

[Getting Started](#)[Training and Certification](#)[AWS Solutions Library](#)[Architecture Center](#)[Product and Technical FAQs](#)[Analyst Reports](#)[AWS Partners](#)

## Developers on AWS

[Developer Center](#)[SDKs & Tools](#)[.NET on AWS](#)[Python on AWS](#)[Java on AWS](#)[PHP on AWS](#)[JavaScript on AWS](#)

[File a Support Ticket](#)[Knowledge Center](#)[AWS re:Post](#)[AWS Support Overview](#)[Legal](#)[AWS Careers](#)[Create an AWS Account](#)

Amazon is an Equal Opportunity Employer: *Minority / Women / Disability / Veteran / Gender Identity / Sexual Orientation / Age.*

## Language

[| عربي](#)[Bahasa Indonesia |](#)[Deutsch |](#)[English |](#)[Español |](#)[Français |](#)[Italiano |](#)[Português |](#)[Tiếng Việt |](#)[Türkçe |](#)[Русский |](#)[ไทย |](#)[日本語 |](#)[한국어 |](#)[中文 \(简体\) |](#)[中文 \(繁體\)](#)

## Privacy

[|](#)[Site Terms](#)[|](#)[Cookie Preferences](#)[|](#)

© 2023, Amazon Web Services, Inc. or its affiliates. All rights reserved.